

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.						
1. REPORT DATE (DD-MM-YYYY) 03-07-2008		2. REPORT TYPE Final Report			3. DATES COVERED (From - To) 03/22/2007 - 03/21/2008	
4. TITLE AND SUBTITLE Adaptive and Robust Multi-Agent Systems					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER N00173-07-1-G010	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Dr. Sean Luke Jeffrey K. Bassett					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) George Mason University 4400 University Drive Fairfax, VA 22030					8. PERFORMING ORGANIZATION REPORT NUMBER 201240	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Ave. SW Washington, DC 20375					10. SPONSOR/MONITOR'S ACRONYM(S) NRL	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT We applied tools based on quantitative genetic theory in order to improve Evolutionary Algorithms for use with team learning tasks. We reviewed the quantitative genetics literature more widely, and developed a theoretical analysis applying genetic theory to the team learning problem. We then constructed and analyzed a neural network structure and new genetic operators which more effectively divide the feature space for the Evolutionary Algorithm. We performed experiments and discovered that the new operators and structure produced more parsimonious results. We plan to publish these results in an upcoming conference following more rigorous experiments.						
15. SUBJECT TERMS Price's Theorem, Quantitative Genetic Theory, Evolutionary Algorithms						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU		3	
					19a. NAME OF RESPONSIBLE PERSON Dr. Sean Luke	
					19b. TELEPHONE NUMBER (Include area code) 703-993-4169	

NRL Grant No. N00173-07-1-G010

Final Report

A Collaborative Research Effort between the Naval Research Laboratory and
George Mason University (2007–2008)

Dr. Sean Luke: Principal Investigator
Jeffrey K. Bassett: Graduate Research Assistant

1 Introduction

The proposal we submitted for this grant was to use tools that we have developed in previous years to improve the robustness and adaptability of agent teams. The assumption was that we would collaborate with Paul Wiegand, who was specifically studying these issues in team learning. Unfortunately, Paul left NRL early during this grant, so we decided to modify our goal.

At a more abstract level, our goal was to use tools based on quantitative genetics [4] theory in order improve Evolutionary Algorithms (EAs) for use with team learning tasks. Up until now our work has focused largely on Price's theorem [6]. We decided to take a step back and review the quantitative genetics literature more widely, with the ultimate goal of applying anything we discovered to team learning.

2 Theory

Price's theorem measures the average effect of selection and variation from one generation to the next. During our previous work [5] [1] [2], we discovered that the variance of these effects is also important to understanding the evolutionary process. We suspected we were not the first to discover this, and decided that a larger review of the quantitative genetics literature was in order.

Our review made it clear that the biologist also place a high emphasis on measuring phenotypic variance. What we noticed though was that all of their equations make assumptions about the form of the genetic representation (DNA of course). One of the reasons that Price's theorem has been so useful for us is that it makes no such assumption. EAs can and do use a wide variety of representations that are not easily analyzed by these equations. LISP trees and Neural Networks are just two such examples. In their current form, these variance equations would only be applicable to a limited number of EAs.

We chose to create a new variance equations that is representation independent. Using Price's equation as a guide, we derived a variance equation which eliminated any reference to genes and alleles, and only measures phenotypic traits. Like Price's equation, the equation is separated into several terms, each of which gives insight into an aspect of the evolutionary process. One term measures the resemblance between parents and offspring (heretability) and the other measures how the phenotype space near the parents are explored (resampling?).

This has helped me to improve my personal model of how an EA works. Much like Eshelman's [3] description of Convergence Controlled Variation, we believe that an EAs heuristic works as follows. An EA searches gradually smaller and smaller sections of the space until a solution is found. At any given time, the area most likely to contain a solution described by the population. Selection is the mechanism by which the scope of the search is reduced, focusing in on more promising areas. The job of the reproductive operators is to resample the same area described by the parent population, thereby continuing the search, but mainly within the promising areas.

3 Application

Using this model of an EA, we was able to reason about issues we had seen in prior work with Pitt Approach rule systems. The custom operators typically used with these EAs can often make large changes to individuals still very late in an evolutionary run. In addition, the representations often used have the can create both over and under specifications in the decision (or feature) space.

Given my new understanding, we reasoned that a representation which was more like a kNN learner's representation would be less likely to result in large phenotypic changes if an individual rule was added or removed. Also, we devised a new crossover operator which performs its cuts in feature space, segregating rules based on which side of a randomly chosen hyperplane they fall, rather than on their position on the genome. we reasoned that this would result in much smaller changes between parents and offspring toward the end of the run.

Experiments performed using these techniques had a very interesting effect. We expected them to result in solutions with improved classification accuracies. Instead the solutions had the same classification accuracy as when more standard representations and operators are used, but the sizes of the rule-sets were much smaller. This approach appears to offer a solution to the long-standing problem of bloat in EA learning. We plan to publish these results in an upcoming conference once we can perform a more rigorous set of experiments.

4 Conclusions

Our collaborative effort this year has produced two main results. First, we have managed to modify quantitative genetics variance equations to a form that will make them applicable the EAs in general. And second, we have used our new understanding of EAs in order to improve an existing EA learning algorithm. These contributions will benefit both NRL and GMU.

There is a clear path here for future work as well. This should include the construction of tools based on the new variance equation to be used for analyzing customized EA representations

and operators. These will be much like way our tools based on Price's theorem. Also, we should use these tools with the Pitt approach algorithm to demonstrate that our modifications really did produce the results we intended.

References

- [1] J. K. Bassett, M. A. Potter, and K. A. De Jong. Looking under the ea hood with price's equation. In *Genetic and Evolutionary Computation Conference – GECCO-2004*. Springer, 2004.
- [2] J. K. Bassett, M. A. Potter, and K. A. De Jong. Applying price's equation to survival selection. In *Proceedings of Genetic and Evolutionary Computation Conference – GECCO-2005*, pages 1371–1378. ACM Press, 2005.
- [3] L. Eshelman, K. Mathias, and J. D. Schaffer. Convergence Controlled Variation. In R. K. Belew and M. D. Vose, editors, *Foundations of Genetic Algorithms 4*, pages 203–224. Morgan Kaufmann, 1996.
- [4] R. A. Falconer. *Introduction to Quantitative Genetics*. Longman, London, 1981.
- [5] M. A. Potter, J. K. Bassett, and K. A. De Jong. Visualizing evolvability with Price's equation. In *Proceedings of the 2003 Congress on Evolutionary Computation*, Canberra, Australia, December 2003. IEEE Press.
- [6] G. Price. Selection and covariance. *Nature*, 227:520–521, 1970.